



(19) Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 193 940 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
03.04.2002 Bulletin 2002/14

(51) Int Cl. 7: H04L 29/06

(21) Application number: 01118847.1

(22) Date of filing: 21.03.1995

(84) Designated Contracting States:
DE GB NL

- Cho, Pong
Cambridge, Massachusetts 02139 (US)
- Chopde, Avinash
Andover, Massachusetts 01810 (US)
- Peters, Eric
Carlisle, Massachusetts 01741 (US)
- Koutavas, Jay
Nashua, New Hampshire 03060 (US)

(30) Priority: 21.03.1994 US 215849

(62) Document number(s) of the earlier application(s) in
accordance with Art. 76 EPC:
95301865.2 / 0 674 414

(71) Applicant: AVID TECHNOLOGY, INC.
Tewksbury, MA 01876 (US)

(74) Representative: Kazi, Ilya et al
Mathys & Squire,
100 Gray's Inn Road
London WC1X 8AL (GB)

(72) Inventors:

- Hamilton, James
Carlisle, Massachusetts 01741 (US)
- Hennessy, Richard
Nashua, New Hampshire 03063 (US)

Remarks:

This application was filed on 14 - 08 - 2001 as a
divisional application to the application mentioned
under TNID code 62.

(54) **Apparatus and computer-implemented process for providing real-time multimedia data transport in a distributed computing system**

(57) The problems of latency, flow control and data loss, and data movement within a client system memory are solved in a distributed multimedia system so as to enable real-time transmission of broadcast quality media data over the network. Latency is reduced by an estimation by the server of client needs. Data loss is pre-

vented and flow control is provided by permitting the server to send only as much information as the network interface can reliably receive. Data movement is minimized by copying data directly from the network interface to memory in a display processor or audio processor.




Fig. 1

Description**FIELD OF THE INVENTION**

[0001] This invention is related to systems and processes for maximizing media data transfer rates in distributed multimedia computer systems.

[0002] More particularly, this invention relates to systems and processes for preventing loss of media data during media data transfer between a client and a server in a distributed multimedia system.

[0003] This invention also more particularly relates to systems and processes for retrieving media data with minimum latency from a disk-based media data server such that a client multimedia application has access to sufficient broadcast quality media data.

BACKGROUND OF THE INVENTION

[0004] A variety of kinds of digital multimedia computer systems are now available. Multimedia systems typically use digital audio and digital image data, especially digital video. It should be understood that in this context, digital video means a sequence of two-dimensional digital still images without any synchronization or other timing information typically associated with a video signal such as a television signal. These systems generally fall into one or more of three broad categories: capture, authoring, and playback.

[0005] Capture involves generating digital media from typically analog media. For example, video signals on a video tape or on film are converted into a sequence of two-dimensional digital still images. These images are typically compressed using known compression systems so as to reduce storage requirements.

[0006] Authoring systems, such as the AVID/1 Media Composer, commercially available from Avid Technology, Inc. of Tewksbury, Massachusetts, allow a user to generate multimedia programs, for example by editing source material or by generating information like computerized graphics. Portions of the source material are called clips. For example, multiple takes of a scene in a film could be captured. An authoring system could be used to edit clips from these takes into an appropriate final version of the scene.

[0007] Playback systems are used to display a multimedia program to a viewer. The program typically is previously edited by an authoring system using previously captured material. The display medium may be, for example, broadcast television or conversion to video signals to create a video tape or video disk, a computer display, etc. Thus, playback generally involves retrieval and transmission of the digital media from storage to the display location according to the specified program and may involve conversion of the digital media to an analog form. Very often an authoring system includes a form of playback system.

[0008] In some industries, all three of these kinds of

systems are used in day-to-day operations. For example, in a broadcast or cable television system, media is captured from satellite, cable or another source. A program is then edited from this material. For example, video and audio footage of news events from around the world may be transmitted via satellite and captured by a broadcast station, which then prepares a news program. The final news program is then played back for broadcast. As another example, in the film industry, media may be captured, and then played back by a director to determine whether any more takes of a given scene would be required. Authoring systems are also becoming more commonly used to facilitate the editing process of a film. After a digital version of the film is completed, this version is used to create a final film product.

[0009] In facilities where there may be many users of the same multimedia material, using multiple types of multimedia systems, a need has arisen to provide these capabilities in a distributed computing environment which allows real-time access to and viewing of multimedia information. Real-time access signifies data rates which support providing broadcast quality video at a client. That is, real-time access involves the ability to play, cut and paste, view a single frame and provide fast forward (or reverse) access to video information at a server. This also may signify an ability to recover from a worst-case disk access. Generally, data transfer rates of about five (5) megabytes per second (Mbps) are preferable. Real-time audio and video systems also require large amounts of data to be transferred. These transfer rates and amounts should be achieved without loss and with minimal latency.

[0010] To this day, no systems are commercially available to handle adequately the data transfer requirements of such multimedia systems, particularly broadcast quality video. Because such distributed systems are not currently commercially-available, users and television stations or film studios or post-production facilities typically have to transport disks containing the multimedia information among different machines in order to perform the various capture, authoring, and playback activities, which is generally undesirable. Alternatively, all of these activities are performed on a single machine.

[0011] For other kinds of applications, there are several conventional methods for transferring large amounts of data between a server and a client in a distributed computer system over a network interconnection. In a first method, a server transmits the data, in bulk, directly to a client. Traditional distributed computing systems utilize a widespread industry practice of a disk server system coupled with an in-memory data cache to provide a mechanism that a client may use to read data. The file system of the underlying operating system performs larger than requested disk reads in an attempt to fill the cache with information that the system estimates the client will need next. One problem with this method is that the client has no control over selection of data, or timing of the transfer.

[0012] In a second method, the client periodically and unpredictably makes requests for data from the server. One problem with this method is that a delay (latency) is incurred while the server prepares the data for delivery, for example, by reading it from a local disk into system memory. This class of systems will not provide the minimum latency required by a media data pipeline operating in a variable consumption rate environment within a video and audio processing system.

[0013] Another problem with current systems is that they have not accounted for resource limitations along the path from the sender to the receiver. That is, when data, such as audio or video media data, is transmitted over a network, a flow control mechanism is required so that the sender does not transmit the data faster than the receiver can accept it. Otherwise, portions of the data packets may be lost in transmission even though the physical interconnect has not failed in any way. This packet loss is normally handled by retransmission, but the real-time requirements of multimedia playback, especially in a broadcast environment, preclude retransmission as a solution.

[0014] The Internet protocol, Transmission Control Protocol/Internet Protocol (TCP/IP), is an example of current networking industry practice. TCP/IP is an example of the use of an end-to-end window extension technique. The receiver extends a window to the sender. The sender sends until this window is consumed. The window only refers to the amount of protocol data unit (PDU) buffer memory available to the application. Network interface level buffering is not considered. If the buffers in the network interface are overflowed or congested, PDUs will be lost. TCP/IP was designed to operate in the face of a PDU loss on the interconnect between the sender and the receiver. The use of timeouts and sequence number checking on each message allows the sender and receiver to determine that a PDU or acknowledge message has been lost or corrupted while moving across the network from such a loss. Transmission retry is used to recover. The overhead of the retry and timeouts reduces the possible data transfer rate below that needed to support processing broadcast quality video and audio information.

[0015] Another common problem in distributed systems is that the receiver typically experiences significant overhead in identifying a final destination in user-level memory into which data is to be copied from an intermediate system buffer and the network interface. This overhead reduces the possible data transfer rate to levels that make impossible the movement of broadcast quality video and audio media data at sufficient speed.

SUMMARY OF THE INVENTION

[0016] The present invention overcomes the problems of latency, flow control and data loss, and minimizes data movement within a client system memory in a distributed multimedia system so as to enable real-time

transmission of broadcast quality media data over the network.

[0017] Latency is reduced by an estimation by the server of client needs. Data loss is prevented and flow control is provided by permitting the server to send only as much information as the network interface can reliably receive. Data movement is minimized by copying data directly from the network interface to memory of input/output devices such as display and audio processors as well as main memory.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] In the drawings,

FIG. 1 is a block diagram of a distributed multi media computing system in accordance with this invention;
 FIG. 2 is a more detailed block diagram of the client and server computers of FIG. 1;
 FIG. 3 is a block diagram describing a media pipeline as is used in the prior art;
 FIG. 4 is a flow chart describing how a server computer prepares to send data to client computer;
 FIG. 5 illustrates a window buffer manager at the server used to anticipate client requests for media data;
 FIG. 6 is an illustration of a packet as may be used in accordance with a protocol in this invention;
 FIG. 7 is a block diagram of a network interface for a client computer; and
 FIG. 8 is a flow chart describing how a client computer receives information.

DETAILED DESCRIPTION

[0019] The present invention will be more completely understood through the following detailed description which should be read in conjunction with the attached drawing in which similar reference numbers indicate similar structures. All references cited herein are hereby expressly incorporated by reference.

[0020] FIG. 1 illustrates a general structure of a distributed multimedia computer system suitable for practicing the invention. It includes a server 40 to which various client systems are interconnected via a high speed media transport 42. The media transport 42 is in a "star" configuration to provide full bandwidth for each client and to support failure recovery. The clients may generally include a capture system 44 and authoring system 46 and a playback system 48 as shown. Many other clients may exist and not all three types of multimedia systems are required.

[0021] To provide low bit error rates, the media transport 42 is preferably optical, e.g. a fiber distributed data interconnect (FDDI) or an asynchronous transfer mode (ATM) system. It should be used only for communicating messages involved in the transfer (i.e., reading and writ-

ing) of media data between the client and the server. There should be no other network traffic on the media transport 42.

[0022] To play two-field broadcast video, each client should have about five (5) megabytes/second (Mbps) of data bandwidth for one video stream and four audio streams. This bandwidth allows for three (3) Mbps of two-field, 100Kb/frame compressed video data, 96Kb/sec for each 48 KHz audio channel and recovery bandwidth. Recovery bandwidth is defined as additional bandwidth needed to recover after a disk head move caused by any of a new video file extent access, a need for audio samples for one or more tracks, or spindle contention from two clients accessing the same disk simultaneously. Thus, with current technology an additional bandwidth of 1.5 Mbps is recommended so that the media transport 42 should allow for five (5) Mbps. Highly scalable bandwidth and single point management and monitoring are also benefits of this kind of interconnect.

[0023] The clients 44, 46, 48 and the server 40 may also be interconnected by a control network 41 which may be, for example, an Ethernet type of network. The control network 41 is used for management of media files and other resources through services of another system unrelated to the communications of media data. The control network 41 is also suitable for general administrative support.

[0024] As with any network, a connection between the server and the client has to be established before sending any data. The connection initialization process is independent of the protocol for media transport. The connection handshaking is typically done over the control network 41. It may also be done using another protocol on the media transport 42.

[0025] As shown in FIG. 2, the server 40 includes a disk 52 on which the media data is primarily stored. The server 40 also includes a processor 54 and system memories 56 which are interconnected with the disk 52 via a bus 58. This bus 58 is also connected to a network interface 60 which connects to the media transport 42. A client 50 includes a central processor 62 and main memory 64 interconnected by a bus 66. Also connected to the bus are preferably video memory 63 and audio memory 65, which stores the multimedia data to be used respectively by display processor 67 and audio processor 69 to display video and play audio. This bus connects to a network interface 68 which connects the client 50 to the media transport 42, and thus to the server 40.

[0026] In the embodiment developed by the inventors, the client computers 44, 46, 48 are Macintosh Quadra 950 computers with twenty-eight (28) megabytes (Mb) of memory and Macintosh System 7.1 system software. The server computer is a Silicon Graphics Challenge L computer with 4 cpu, 100 megahertz (Mhz), 256 Mb of memory, 8 gigabytes (gb) of disk storage and the IRIX 5.2 operating system. The media transport 42 and network interface 60 is provided by an optical fiber with the FORE Systems ForeRunner VMA-200 ATM VMEbus

Adapter for the server 40, and associated system software, and application program interface (API) library software. For the clients, a FORE Systems ForeRunner NBA-200 ATM NUBUS Adapter was used as the network interface 68, along with associated system software, and API library software. It should be understood that the invention is not limited to these specific machines and components mentioned and that other network interfaces and computers may be used.

[0027] The need for reading media data from the server 40 into a client is dependent on the need for the media data further down the processing stream, particularly by the display processor and audio processor. That is, this media data should be considered as being used in a sort of media pipeline, such as that described in U.S. Patent No. 5,045,940, issued September 3, 1991 and assigned to Avid Technology, Inc of Tewksbury, Massachusetts. In this system, as shown simplified in FIG. 3, video data is input through an input/output system 100. An audio stream is not shown. The data then goes to a frame buffer 102 from where it is then compressed using a decompression/compression system 104. Compressed data is stored in a buffer 106 from where it is then transferred into permanent storage 108. Conversely, data can be read from permanent storage 108 to the compressed data buffer 106 from where it is decompressed in 104 and applied to a frame buffer 102 from where it is then displayed by an output unit 100. In this system, processing by any one of the subsystems is delayed if the subsystem to which its output is provided cannot receive any more data. The general concept is to maintain the pipeline at a completely full state in all of its stages so as to maximize the quality of the video being seen. The distributed system in this invention would be primarily used so as to provide permanent storage 108 at a server 40 (FIG. 1) and to have the remainder of the subsystems 100-106 at a client. Thus, requests for media data from the server 40 come at times indicated by other subsystems of a media pipeline as media data is needed.

[0028] At the application level, there are two types of messages in this system: protocol requests/responses and media data packets. Protocol requests/responses occur at the application level of the interconnect between the client and the server. Media data transfer occurs at the transport level. Both types of messages are communicated over the media transport 42.

[0029] Having described a general structure of a distributed computer system and a network protocol for supporting this invention, its use in providing real-time broadcast quality video will now be described. Three general concepts are used to obtain the desired results. First of all, clients send some information to the server 40 in advance of requests for media data, which information allows the server to access the data from the disk and place it in faster system memory in advance of client requests, thus reducing latency. The server, in effect, is allowed to predict the requests from the client. This information is sent by the client to the server as a protocol

request/response at the application level. Lossless flow control is provided by having a client authorize the server 40 to send only an amount of data which the network interface is capable of handling. Finally, protocol control and identity mechanisms are provided within data packets to be transferred. These protocol control and identity mechanisms are then used by the network interface to process the data and automatically and directly transmit it to an appropriate user-level memory buffer to avoid copying of data into system level memory at the client. These concepts will now be described in more detail.

[0030] How the server is informed in advance of the needs of the client will first be described. In either an authoring system or a playback system, it is simple to predict the media data which will be required by a client. For example, in a typical authoring system, the user selects from a directory a set of clips which will be used in the product to be edited. Upon the selection of these clips, a message may be sent to the server by the client requesting that those clips be prepared for delivery. Similarly, in a playback system, one typically knows in advance all of the clips which will be used in the program to be played. Both of these types of prediction can be readily performed using the Avid/1 Media Composer. For example, when editing using the Avid/1 Media Composer, a user selects a number of clips by selection from a directory. Also, as a program is being edited, a data structure is developed called a composition which indicates the relationships and ordering of clips used in the program. A list of clips needed can be readily determined and sent to the server to permit the server to prepare in advance the data needed to display these clips.

[0031] The sequence of steps performed to effect this prediction will now be discussed in connection with FIG. 4. First, the list of media objects to be used by the client is created by the client in step 120. The client then requests that the server create a data structure, herein called a TrackList, for storing the media objects for the list in step 122. The TrackList is a resource on the server. Subsequent requests by the client are used to populate, or define the contents of, the TrackList. The order of these requests dictates the order of the entries in the TrackList, as will be described below. The server responds to the client with an identifier for the TrackList which the client then includes in any requests to the server for associating media objects with the TrackList. The server may refuse a request to create a TrackList due to lack of resources and reply with an error in step 124.

[0032] After an identifier of a TrackList is received, the client then sends, in step 126 a series of requests to open media files for media objects corresponding to the list created by the client. Each request includes an identifier for the media object to be opened, the identifier of the TrackList, and an indication of whether the media object is to be inserted at the beginning or end of the TrackList. The identifier of the media object includes the media file and a range of bytes within that media file.

[0033] For each request to obtain a media object, the server opens the media object in step 128 and associates it with the TrackList. The operation performed by the server opens a logical read stream to an existing

5 media file, hereinafter called a flow, located on the server. The flow is a cache buffer which is allocated and linked to the TrackList. Thus, the TrackList is a list of flows. It is possible to obtain single frame access to media files by specifying a NULL identifier for the TrackList.

10 In this instance, no resource reservations are made.

[0034] If this server contains the media file identified by the client request, it returns an identifier to the flow that was opened by it in step 130, which identifier the client includes with subsequent requests to read media 15 data from this file. If a requested media file is not located on the server an error is returned. If an incorrect identifier for the TrackList is provided, an error is also returned. If the media file requested is in use by another client, an error is also issued due to a resource conflict.

20 [0035] After the appropriate requests to open media files for read access, the media data desired by a client for a track are placed in a window buffer (described below) associated with that track so that the media data are ready to be accessed upon a read request. Using 25 the predictive information in the TrackList the server will already have read the needed media data into the window buffer and can immediately fulfill the request.

[0036] As shown in FIG. 5, a window manager 140 at the server 40 is responsible for maintaining a window 30 about a currently active pointer 144 in the window buffer 142 by doing autonomous read forward or backward operations. Reading sample data with minimum server turn-around delay is thus obtained because the server can do look ahead disk reads to fill the window buffer

35 142 for the next flow while the current flow is still being read by client. The TrackList allows the server to autonomously look ahead according to the information about the order of play of the streams, thus maintaining this window buffer 142 full. These read ahead operations for the media objects referenced by the TrackList can be performed such that the window buffer 142 will always contain the exact data needed to satisfy subsequent read operations.

40 [0037] When the media data is prepared in the window buffer by the server, the server 40 is ready to handle read requests from the client, which will now be described in connection with Figs. 6-10. The preparation activities are performed at the application level; mechanisms for reading are performed at the transport level 45 as will be described below.

[0038] To enhance throughput for reading media data in this system, a new network interface and protocol was developed for use at the transport level.

50 [0039] As shown in FIG. 7, the network interface 68 includes a memory 80, preferably a first-in, first-out (FIFO) memory for receiving packets from the media transport 42. The memory 80 is controlled by a controller 82 which reads packets from the memory 80, and then

transfers the data portion of those packets, thus demultiplexing the packets, to logically separate direct memory access (DMA) elements 84, 86 and 88 which are connected to the system bus 66. Media data of a video track goes to the logical DMA to video memory 66. Similarly, audio data is sent to the logical DMA to audio memory 88. Finally, data concerning protocol requests/responses and window packets are sent to the logical DMA to main memory 84.

[0040] As discussed above, the client creates a TrackList for each track in the media program to be played, authored or captured. Likewise, a receive channel is also created by the client for each of these tracks to receive media data. The server similarly opens the same number of transmit channels, each associated with a TrackList. A channel is a logical construct used to control flow of data between the network interface and main memory or audio or video memory. More particularly, a channel is a half-duplex virtual circuit with a logical DMA control scheme, a flow control scheme to be described below and associated state information as discussed below. Each channel has a separate logical DMA controller. Thus the interface of FIG. 7 represents two media data channels. Both the server and the client have a similar command channel which is used for communicating data concerning protocol requests and responses. Each transmit channel at the server is organized as a cache, or a window buffer, into the stream of bytes being read by the client as discussed above in connection with FIG. 5.

[0041] The network interface maintains a set of state information for each active receive channel. This information includes 1) an offset in bytes from the start of a packet to the start of the field storing the identifier for the receive channel, 2) an indication of the length in bytes of the identifier of this receive channel, 3) the identifier used by the network interface to identify this receive channel, 4) an offset in bytes from the start of a packet to the start of data, 5) a starting address of the destination memory to which the network interface will move the data, 6) a queue of all receive requests for this channel and including the address and size of all transfer units to be sent authorization indicating the number and size of all transfer units authorized to be sent to this receive channel (to be described below), and 7) a count in bytes of data transferred by the network interface for this receive channel to this point in time. For each channel items 1), 2) and 4) are long lasting and should not change in real time. The network interface maintains state variables for the remainder. The network interface also maintains a global state variable indicative of an amount of data that it has authorized to be received.

[0042] The network interface also maintains a set of state information for each active transmit channel. This information includes 1) an offset in bytes from the start of a packet to the start of the identifier for a channel, 2) an indication of the length in bytes of the identifier, 3) an indication of the token used by the network interface to

identify the transmit channel, 4) an offset in bytes from the start of the packet to the start of data, 5) a starting address of the source buffer from which the network interface will move the data, 6) an offset in the source buffer from where the media data will be taken, 7) an indication of the total number of data bytes to transfer from the indicated address, and 8) stream protocol header information, used by the network interface to encapsulate the retrieved data bytes into a proper network packet.

5 The network interface also keeps track of the number and size of transfer units authorized for each transmit channel and the number of outstanding bytes authorized to be transmitted to a given destination address.

[0043] This new network interface was defined by using a protocol handler which is part of a device driver for the network interface. The protocol handler demultiplexes received packets on the basis of protocol type. An indication of the protocol type of a packet, often called a discriminator, is located in a predetermined location in the packet. A typical packet for this new protocol is shown in FIG. 4.

[0044] In this protocol, as shown in FIG. 4, media data is encapsulated within a packet 70. This format includes a network header 72, a protocol header 74, data 76, and a packet or package check sum 78. The network header 72 includes the indications of the sender and receiver of the packet, as well as the discriminator indicating the protocol type. The protocol header indicates the type of packet within the protocol being received. At the transport level, there are two types of packets for media transport in this protocol: data packets and window packets.

[0045] Data packets contain an arbitrary amount of data. Such packets may be media data packets for media channels or protocol request/responses data for command channels. A data packet contains in the protocol header 74 an indication of the receive channel at the destination and, in data portion 76, a block of data destined for the receive channel indicated. The amount of data in the packet and the offset in bytes in the destination buffer where the network interface will move the data are also included in the protocol header 74.

[0046] Window packets indicate information used by the server to control the amount of data sent to the client. A window packet includes an indication of the channel at the source of the window packet to which data is to be sent (i.e., the destination of the requested data) an indication of whether an additional transfer unit is authorized, as well as a number of bytes which is authorized to be sent. A window packet authorizes the server to transmit either one additional transfer unit on a given virtual channel or N additional bytes on the network interface, or both. The first of these allows the source of the window packet to prevent the transmission of packets for which no receive buffers are available. The second ensures that the flow of data does not overrun the capabilities of the network interface card to buffer packets for delivery to the source of the window packet. How these window packets are used will be described in

more detail below.

[0047] The network interface (FIG. 7) is used in the following manner to control flow between the sender or server 40 and client 50. First, as mentioned above for each track, a virtual receive channel is created. That is, for example, a video program typically has at least one track of video data and one track of audio data. A virtual receive channel is created for each of these tracks by the client in step 90 of FIG. 8. For each virtual receive channel at the client a logical DMA is created in the network interface 68. When the client places a request to the server for media data, an identifier of the receive channel to which media data will be sent is included, along with an indication of the media file to be accessed and an amount of data requested from it, called a transfer unit, is also provided. This request is made at the application level after the TrackList for the channel is established. Next, at the transport level, a window packet is sent in step 92, authorizing one transfer unit to be sent to this receive channel. However, the virtual receive channel only has a given size S, i.e., a maximum amount of data that it may hold. Thus, the window packet authorizing a transfer unit only authorizes a number of bytes to be sent which is less than the size S of the virtual channel. Typically, the first transfer for a program includes both, as identified above. However, if an application is playing a number of media files, a media file in the middle of the program may be requested while another is already in memory at the client. Thus, a transfer unit may be authorized by such a request, but no data in addition to that authorized previously may be sent. This method assumes that there is no other traffic on the network interface that is under its control. It also assumes the existence of a network interface 68 that is capable of reliably buffering at least one complete data packet at the full speed of the network.

[0048] The network interface processes a transmit channel by reading network packet size data chunks from the specified addresses, with the last chunk being of arbitrary length, while encapsulating the data in proper header information then sending completed packets to the destination node. The network interface only reads as many bytes as specified. The data packet sent also includes the identifier of the receive channel previously provided by the client.

[0049] The received data is then processed by the client in step 94.

[0050] For reception, the data must be separated from the packet and moved to buffers for use. Traditionally this separation is performed by the client after the packet has been moved into system memory by the network interface. The client then moves the data to the appropriate destination using a software copy loop. However, no system memory buffer copy is possible given the data rates required of broadcast quality video. This traditional method is bypassed in this invention to enhance throughput by minimizing data copying using direct memory access. The identifier of the receive channel is

used by the network interface under direct memory access (DMA) control to directly transfer the media data from the network interface into the system memory or other I/O device memory, such as video memory 63 or audio memory 65, thus enhancing throughput, as shown in FIG. 7. After the received data is processed, a window authorization is sent again for more data, in step 92. Thus, a series of network packets is received into the network interface are decomposed from the packets then written directly from the network interface to specific destination memory addresses with media data processing hardware. The network interface uses the state variables of the receive data channel to move the data portion of the packet into the destination addresses.

[0051] As an example, assume that a client wishes to have one megabyte of information from a media file. A read request is sent from the client to the server at the application level to request that amount of data. The one megabyte of information represents a transfer unit. Next, a window authorization is then sent to the server at the transport level authorizing one transfer unit to be sent, but that only 16 kilobytes of data to be sent, assuming that the size of the packet FIFO (80 in FIG. 7) at the client is 16 kilobytes. The server then sends the appropriate data from these senders transmit stream only sending a series of packets up to 16 kilobytes. One packet is likely not to be a total of 16 kilobytes, it may be only a small number, e.g., typically approximately 4K bytes. As the client processes the received data, at some point, another window authorization may be sent, for example, authorizing an additional 8K of bytes to be sent. The client will do this after a certain number of bytes of buffer memory in the network interface become available. The server continues to send media data until the number of bytes authorized is exhausted or until the requested transfer unit is exhausted.

[0052] Aside from providing real-time access to broadcast quality media data, this system provides numerous advantages over standard multimedia systems. For example, recording digitized media directly onto the server allows a user at a client computer to play video without having to copy large media files from the server to the client computer. Also, local disk space at each client computer is not needed to hold copies of media files. One copy of each file at the server is all that is needed.

[0053] Having now described a few embodiments of the invention, it should be apparent to those skilled in the art that the foregoing is merely illustrative and not limiting, having been presented by way of example only. Numerous modifications and other embodiments are within the scope of one of ordinary skill in the art. For example, fault tolerance may be provided in this system by providing redundancy, for example, redundant interconnect and network subsystems can be added as desired to improve system robustness. These and other modifications are contemplated as falling within the

scope of the invention as defined by the appended claims.

[0054] The invention may also be defined by the following numbered clauses.

1. A distributed multimedia network system, comprising:

a server computer having a disk on which media data is stored;
 a client computer having means for requesting media data from the server computer and means for providing to the server computer a list of media data which may be requested by the client;
 a media transport connecting the server computer to the client computer and communicating media data between the client computer and the server computer; and
 in the server computer, means for retrieving media data from the disk as indicated in the list provided by the client computer and for caching the retrieved media data.

2. The distributed multimedia network system of clause 1,

wherein the server computer includes means for sending to the client computer an identifier which references the cached media data at the server,
 wherein the client computer further comprises means for sending requests to the server, with the identifier, for a next portion of the media data, and
 wherein the server computer, in response to said requests, sends the requested media data to the client computer.

3. The distributed multimedia network system of clause 2,

wherein the means for retrieving, in the server computer, only retrieves a first portion of the media data in the list at a given time, and wherein the means for retrieving, in the server computer, retrieves a latter portion of the media data in the list after the client computer requests the first portion of the media data.

4. The distributed multimedia network system of clause 2,

wherein the requests from the client include an indication of an amount of data to be transferred, and
 wherein the client computer includes means for determining the amount of data to be trans-

ferred according to capacity of a network interface in the client computer.

5. The distributed multimedia network system of clause 1,

wherein the client computer includes a network interface,
 wherein the server computer includes means for sending cached media data to the client computer including an indication of an application at the client computer to receive the media data, and
 wherein the network interface of the client computer includes means for receiving the media data from the server computer and for transmitting the media data directly to memory in the client computer used by the application without copying the media data into system level memory in the client computer.

6. The distributed multimedia network system of clause 1,

wherein the server computer maintains a window buffer in which media data is cached, and wherein the means for retrieving and for caching retrieves and caches only enough media data to fill the window buffer.

7. The distributed multimedia network system of clause 6,

wherein the server computer, in response to requests from the client, includes means for sending media data in the window buffer to the client computer, and
 wherein the server computer includes means for operating the means for retrieving and caching so as to fill the window buffer after sending media data from the window buffer to the client computer.

8. The distributed multimedia network system of clause 1, wherein the media data is full-motion broadcast quality video.

9. The distributed multimedia network system of clause 8, wherein the media data also includes audio associated with the video.

Claims

55 1. A process for directing video and audio data from a computer network to output devices, comprising:

associating a video output device having a

memory with a first channel having a direct memory access engine;
 associating an audio output device having a memory with a second channel having a direct memory access engine;
 receiving data packets from a computer network into a buffer, wherein each data packet includes header information and data, including an indication of one of the first and second channels; and
 for each data packet, transferring the data in the data packet directly from the buffer to a memory associated with the output device corresponding to the indication of the first and second channels using the direct memory access engine of the channel, thereby bypassing copying of the data by the operating system to system level memory.

2. A network interface for directing video and audio data from a computer network to output devices, comprising:
 means for associating a video output device having a memory with a first channel having a direct memory access engine;
 means for associating an audio output device having a memory with a second channel having a direct memory access engine;
 means for receiving data packets from a computer network into a buffer, wherein each data packet includes header information and data, including an indication of one of the first and second channels; and
 means for transferring the data in each data packet directly from the buffer to the memory of the output device corresponding to the indication of the first and second channels using the direct memory access engine of the channel, thereby bypassing copying of the data by the operating system to system level memory.

3. A network interface for directing video and audio data from a computer network to output devices, comprising:
 a first channel having a direct memory access engine for accessing memory in a video output device;
 a second channel having a direct memory access engine for accessing memory in an audio output device;
 a buffer for receiving data packets from a computer network, wherein each data packet includes header information and data, including an indication of one of the first and second channels; and
 a controller for transferring the data in each data packet directly from the buffer to the memory of the output device corresponding to the indication of the first and second channels using the direct memory access engine of the channel, thereby bypassing copying of the data by the operating system to system level memory.

5 4. A process for controlling flow of data over a computer network from a server to a client, comprising:
 sending a request to a server indicating a channel at the client, a file to be accessed and an amount of data to be read from the file;
 sending a message to a server authorizing data to be transferred from the file in an amount less than an amount of space available in a buffer on the network interface;
 receiving data packets from the server and storing the data packets into the buffer, wherein each data packet includes an indication of the channel at the client;
 processing the received data to provide the received data from the buffer to a device associated with the indicated channel; and
 repeating steps of sending the message, receiving data packets and processing the received data until the amount of data to be read from the file is received.

10 5. The process of claim 4, wherein the data comprises video and audio data to be played back on respective output devices, wherein the process further comprises:
 associating a video output device having a memory with a first channel having a direct memory access engine;
 associating an audio output device having a memory with a second channel having a direct memory access engine;

15 40 45 50 55 60 65 70 75 80 85 90 95 100 105 110 115 120 125 130 135 140 145 150 155 160 165 170 175 180 185 190 195 200 205 210 215 220 225 230 235 240 245 250 255 260 265 270 275 280 285 290 295 300 305 310 315 320 325 330 335 340 345 350 355 360 365 370 375 380 385 390 395 400 405 410 415 420 425 430 435 440 445 450 455 460 465 470 475 480 485 490 495 500 505 510 515 520 525 530 535 540 545 550 555 560 565 570 575 580 585 590 595 600 605 610 615 620 625 630 635 640 645 650 655 660 665 670 675 680 685 690 695 700 705 710 715 720 725 730 735 740 745 750 755 760 765 770 775 780 785 790 795 800 805 810 815 820 825 830 835 840 845 850 855 860 865 870 875 880 885 890 895 900 905 910 915 920 925 930 935 940 945 950 955 960 965 970 975 980 985 990 995 1000 1005 1010 1015 1020 1025 1030 1035 1040 1045 1050 1055 1060 1065 1070 1075 1080 1085 1090 1095 1100 1105 1110 1115 1120 1125 1130 1135 1140 1145 1150 1155 1160 1165 1170 1175 1180 1185 1190 1195 1200 1205 1210 1215 1220 1225 1230 1235 1240 1245 1250 1255 1260 1265 1270 1275 1280 1285 1290 1295 1300 1305 1310 1315 1320 1325 1330 1335 1340 1345 1350 1355 1360 1365 1370 1375 1380 1385 1390 1395 1400 1405 1410 1415 1420 1425 1430 1435 1440 1445 1450 1455 1460 1465 1470 1475 1480 1485 1490 1495 1500 1505 1510 1515 1520 1525 1530 1535 1540 1545 1550 1555 1560 1565 1570 1575 1580 1585 1590 1595 1600 1605 1610 1615 1620 1625 1630 1635 1640 1645 1650 1655 1660 1665 1670 1675 1680 1685 1690 1695 1700 1705 1710 1715 1720 1725 1730 1735 1740 1745 1750 1755 1760 1765 1770 1775 1780 1785 1790 1795 1800 1805 1810 1815 1820 1825 1830 1835 1840 1845 1850 1855 1860 1865 1870 1875 1880 1885 1890 1895 1900 1905 1910 1915 1920 1925 1930 1935 1940 1945 1950 1955 1960 1965 1970 1975 1980 1985 1990 1995 2000 2005 2010 2015 2020 2025 2030 2035 2040 2045 2050 2055 2060 2065 2070 2075 2080 2085 2090 2095 2100 2105 2110 2115 2120 2125 2130 2135 2140 2145 2150 2155 2160 2165 2170 2175 2180 2185 2190 2195 2200 2205 2210 2215 2220 2225 2230 2235 2240 2245 2250 2255 2260 2265 2270 2275 2280 2285 2290 2295 2300 2305 2310 2315 2320 2325 2330 2335 2340 2345 2350 2355 2360 2365 2370 2375 2380 2385 2390 2395 2400 2405 2410 2415 2420 2425 2430 2435 2440 2445 2450 2455 2460 2465 2470 2475 2480 2485 2490 2495 2500 2505 2510 2515 2520 2525 2530 2535 2540 2545 2550 2555 2560 2565 2570 2575 2580 2585 2590 2595 2600 2605 2610 2615 2620 2625 2630 2635 2640 2645 2650 2655 2660 2665 2670 2675 2680 2685 2690 2695 2700 2705 2710 2715 2720 2725 2730 2735 2740 2745 2750 2755 2760 2765 2770 2775 2780 2785 2790 2795 2800 2805 2810 2815 2820 2825 2830 2835 2840 2845 2850 2855 2860 2865 2870 2875 2880 2885 2890 2895 2900 2905 2910 2915 2920 2925 2930 2935 2940 2945 2950 2955 2960 2965 2970 2975 2980 2985 2990 2995 3000 3005 3010 3015 3020 3025 3030 3035 3040 3045 3050 3055 3060 3065 3070 3075 3080 3085 3090 3095 3100 3105 3110 3115 3120 3125 3130 3135 3140 3145 3150 3155 3160 3165 3170 3175 3180 3185 3190 3195 3200 3205 3210 3215 3220 3225 3230 3235 3240 3245 3250 3255 3260 3265 3270 3275 3280 3285 3290 3295 3300 3305 3310 3315 3320 3325 3330 3335 3340 3345 3350 3355 3360 3365 3370 3375 3380 3385 3390 3395 3400 3405 3410 3415 3420 3425 3430 3435 3440 3445 3450 3455 3460 3465 3470 3475 3480 3485 3490 3495 3500 3505 3510 3515 3520 3525 3530 3535 3540 3545 3550 3555 3560 3565 3570 3575 3580 3585 3590 3595 3600 3605 3610 3615 3620 3625 3630 3635 3640 3645 3650 3655 3660 3665 3670 3675 3680 3685 3690 3695 3700 3705 3710 3715 3720 3725 3730 3735 3740 3745 3750 3755 3760 3765 3770 3775 3780 3785 3790 3795 3800 3805 3810 3815 3820 3825 3830 3835 3840 3845 3850 3855 3860 3865 3870 3875 3880 3885 3890 3895 3900 3905 3910 3915 3920 3925 3930 3935 3940 3945 3950 3955 3960 3965 3970 3975 3980 3985 3990 3995 4000 4005 4010 4015 4020 4025 4030 4035 4040 4045 4050 4055 4060 4065 4070 4075 4080 4085 4090 4095 4100 4105 4110 4115 4120 4125 4130 4135 4140 4145 4150 4155 4160 4165 4170 4175 4180 4185 4190 4195 4200 4205 4210 4215 4220 4225 4230 4235 4240 4245 4250 4255 4260 4265 4270 4275 4280 4285 4290 4295 4300 4305 4310 4315 4320 4325 4330 4335 4340 4345 4350 4355 4360 4365 4370 4375 4380 4385 4390 4395 4400 4405 4410 4415 4420 4425 4430 4435 4440 4445 4450 4455 4460 4465 4470 4475 4480 4485 4490 4495 4500 4505 4510 4515 4520 4525 4530 4535 4540 4545 4550 4555 4560 4565 4570 4575 4580 4585 4590 4595 4600 4605 4610 4615 4620 4625 4630 4635 4640 4645 4650 4655 4660 4665 4670 4675 4680 4685 4690 4695 4700 4705 4710 4715 4720 4725 4730 4735 4740 4745 4750 4755 4760 4765 4770 4775 4780 4785 4790 4795 4800 4805 4810 4815 4820 4825 4830 4835 4840 4845 4850 4855 4860 4865 4870 4875 4880 4885 4890 4895 4900 4905 4910 4915 4920 4925 4930 4935 4940 4945 4950 4955 4960 4965 4970 4975 4980 4985 4990 4995 5000 5005 5010 5015 5020 5025 5030 5035 5040 5045 5050 5055 5060 5065 5070 5075 5080 5085 5090 5095 5100 5105 5110 5115 5120 5125 5130 5135 5140 5145 5150 5155 5160 5165 5170 5175 5180 5185 5190 5195 5200 5205 5210 5215 5220 5225 5230 5235 5240 5245 5250 5255 5260 5265 5270 5275 5280 5285 5290 5295 5300 5305 5310 5315 5320 5325 5330 5335 5340 5345 5350 5355 5360 5365 5370 5375 5380 5385 5390 5395 5400 5405 5410 5415 5420 5425 5430 5435 5440 5445 5450 5455 5460 5465 5470 5475 5480 5485 5490 5495 5500 5505 5510 5515 5520 5525 5530 5535 5540 5545 5550 5555 5560 5565 5570 5575 5580 5585 5590 5595 5600 5605 5610 5615 5620 5625 5630 5635 5640 5645 5650 5655 5660 5665 5670 5675 5680 5685 5690 5695 5700 5705 5710 5715 5720 5725 5730 5735 5740 5745 5750 5755 5760 5765 5770 5775 5780 5785 5790 5795 5800 5805 5810 5815 5820 5825 5830 5835 5840 5845 5850 5855 5860 5865 5870 5875 5880 5885 5890 5895 5900 5905 5910 5915 5920 5925 5930 5935 5940 5945 5950 5955 5960 5965 5970 5975 5980 5985 5990 5995 6000 6005 6010 6015 6020 6025 6030 6035 6040 6045 6050 6055 6060 6065 6070 6075 6080 6085 6090 6095 6100 6105 6110 6115 6120 6125 6130 6135 6140 6145 6150 6155 6160 6165 6170 6175 6180 6185 6190 6195 6200 6205 6210 6215 6220 6225 6230 6235 6240 6245 6250 6255 6260 6265 6270 6275 6280 6285 6290 6295 6300 6305 6310 6315 6320 6325 6330 6335 6340 6345 6350 6355 6360 6365 6370 6375 6380 6385 6390 6395 6400 6405 6410 6415 6420 6425 6430 6435 6440 6445 6450 6455 6460 6465 6470 6475 6480 6485 6490 6495 6500 6505 6510 6515 6520 6525 6530 6535 6540 6545 6550 6555 6560 6565 6570 6575 6580 6585 6590 6595 6600 6605 6610 6615 6620 6625 6630 6635 6640 6645 6650 6655 6660 6665 6670 6675 6680 6685 6690 6695 6700 6705 6710 6715 6720 6725 6730 6735 6740 6745 6750 6755 6760 6765 6770 6775 6780 6785 6790 6795 6800 6805 6810 6815 6820 6825 6830 6835 6840 6845 6850 6855 6860 6865 6870 6875 6880 6885 6890 6895 6900 6905 6910 6915 6920 6925 6930 6935 6940 6945 6950 6955 6960 6965 6970 6975 6980 6985 6990 6995 7000 7005 7010 7015 7020 7025 7030 7035 7040 7045 7050 7055 7060 7065 7070 7075 7080 7085 7090 7095 7100 7105 7110 7115 7120 7125 7130 7135 7140 7145 7150 7155 7160 7165 7170 7175 7180 7185 7190 7195 7200 7205 7210 7215 7220 7225 7230 7235 7240 7245 7250 7255 7260 7265 7270 7275 7280 7285 7290 7295 7300 7305 7310 7315 7320 7325 7330 7335 7340 7345 7350 7355 7360 7365 7370 7375 7380 7385 7390 7395 7400 7405 7410 7415 7420 7425 7430 7435 7440 7445 7450 7455 7460 7465 7470 7475 7480 7485 7490 7495 7500 7505 7510 7515 7520 7525 7530 7535 7540 7545 7550 7555 7560 7565 7570 7575 7580 7585 7590 7595 7600 7605 7610 7615 7620 7625 7630 7635 7640 7645 7650 7655 7660 7665 7670 7675 7680 7685 7690 7695 7700 7705 7710 7715 7720 7725 7730 7735 7740 7745 7750 7755 7760 7765 7770 7775 7780 7785 7790 7795 7800 7805 7810 7815 7820 7825 7830 7835 7840 7845 7850 7855 7860 7865 7870 7875 7880 7885 7890 7895 7900 7905 7910 7915 7920 7925 7930 7935 7940 7945 7950 7955 7960 7965 7970 7975 7980 7985 7990 7995 8000 8005 8010 8015 8020 8025 8030 8035 8040 8045 8050 8055 8060 8065 8070 8075 8080 8085 8090 8095 8100 8105 8110 8115 8120 8125 8130 8135 8140 8145 8150 8155 8160 8165 8170 8175 8180 8185 8190 8195 8200 8205 8210 8215 8220 8225 8230 8235 8240 8245 8250 8255 8260 8265 8270 8275 8280 8285 8290 8295 8300 8305 8310 8315 8320 8325 8330 8335 8340 8345 8350 8355 8360 8365 8370 8375 8380 8385 8390 8395 8400 8405 8410 8415 8420 8425 8430 8435 8440 8445 8450 8455 8460 8465 8470 8475 8480 8485 8490 8495 8500 8505 8510 8515 8520 8525 8530 8535 8540 8545 8550 8555 8560 8565 8570 8575 8580 8585 8590 8595 8600 8605 8610 8615 8620 8625 8630 8635 8640 8645 8650 8655 8660 8665 8670 8675 8680 8685 8690 8695 8700 8705 8710 8715 8720 8725 8730 8735 8740 8745 8750 8755 8760 8765 8770 8775 8780 8785 8790 8795 8800 8805 8810 8815 8820 8825 8830 8835 8840 8845 8850 8855 8860 8865 8870 8875 8880 8885 8890 8895 8900 8905 8910 8915 8920 8925 8930 8935 8940 8945 8950 8955 8960 8965 8970 8975 8980 8985 8990 8995 9000 9005 9010 9015 9020 9025 9030 9035 9040 9045 9050 9055 9060 9065 9070 9075 9080 9085 9090 9095 9100 9105 9110 9115 9120 9125 9130 9135 9140 9145 9150 9155 9160 9165 9170 9175 9180 9185 9190 9195 9200 9205 9210 9215 9220 9225 9230 9235 9240 9245 9250 9255 9260 9265 9270 9275 9280 9285 9290 9295 9300 9305 9310 9315 9320 9325 9330 9335 9340 9345 9350 9355 9360 9365 9370 9375 9380 9385 9390 9395 9400 9405 9410 9415 9420 9425 9430 9435 9440 9445 9450 9455 9460 9465 9470 9475 9480 9485 9490 9495 9500 9505 9510 9515 9520 9525 9530 9535 9540 9545 9550 9555 9560 9565 9570 9575 9580 9585 9590 9595 9600 9605 9610 9615 9620 9625 9630 9635 9640 9645 9650 9655 9660 9665 9670 9675 9680 9685 9690 9695 9700 9705 9710 9715 9720 9725 9730 9735 9740 9745 9750 9755 9760 9765 9770 9775 9780 9785 9790 9795 9800 9805 9810 9815 9820 9825 9830 9835 9840 9845 9850 9855 9860 9865 9870 9875 9880 9885 9890 9895 9900 9905 9910 9915 9920 9925 9930 9935 9940 9945 9950 9955 9960 9965 9970 9975 9980 9985 9990 9995 10000 10005 10010 10015 10020 10025 10030 10035 10040 10045 10050 10055 10060 10065 10070 10075 10080 10085 10090 10095 10100 10105 10110 10115 10120 10125 10130 10135 10140 10145 10150 10155 10160 10165 10170 10175 10180 10185 10190 10195 10200 10205 10210 10215 10220 10225 10230 10235 10240 10245 10250 10255 10260 10265 10270 10275 10280 10285 10290 10295 10300 10305 10310 10315 10320 10325 10330 10335 10340 10345 10350 10355 10360 10365 10370 10375 10380 10385 10390 10395 10400 10405 10410 10415 10420 10425 10430 10435 10440 10445 10450 10455 10460 10465 10470 10475 10480 10485 10490 10495 10500 10505 10510 10515 10520 10525 10530 10535 10540 10545 10550 10555 10560 10565 10570 10575 10580 10585 10590 10595 10600 10605 10610 10615 10620 10625 10630 10635 10640 10645 10650 10655 10660 10665 10670 10675 10680 10685 10690 10695 10700 10705 10710 10715 10720 10725 10730 10735 10740 10745 10750 10755 10760 10765 10770 10775 10780 10785 10790 10795 10800 10805 10810 10815 10820 10825 10830 10835 10840 10845 10850 10855 10860 10865 10870 10875 10880 10885 10890 10895 10900 10905 10910 10915 10920 10925 10930 10935 10940 10945 10950 10955 10960 10965 10970 10975 10980 10985 10990 10995 11000 11005 11010 11015 11020 11025 11030 11035 11040 11045 11050 11055 11060 11065 11070 11075 11080 11085 11090 11

source allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read an amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

7. The process of claim 4, further comprising:

sending the server a list of files which may be accessed by the client;

receiving from the server an indicator of a resource allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read an amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

8. An apparatus for controlling flow of data over a computer network from a server to a client, comprising:

means for sending a request to a server indicating a channel at the client, a file to be accessed and a desired amount of data to be read from the file;

means for sending a message to a server authorizing data to be transferred from the file in an amount less than an amount of space available in a buffer on the network interface;

means for receiving data packets from the server and storing the data packets into the buffer, wherein each data packet includes an indication of the channel at the client;

means for processing the received data to provide the received data from the buffer to a device associated with the indicated channel; and

wherein the means for sending the message, means for receiving data packets and means for processing the received data cooperate to cause the desired amount of data to be received from the file.

9. The apparatus of claim 8, wherein the data comprises video and audio data to be played back on respective output devices, wherein the apparatus further comprises:

means for associating a video output device having a memory with a first channel having a direct memory access engine;

means for associating an audio output device having a memory with a second channel having a direct memory access engine;

wherein the indication of the channel includes an indication of one of the first and second channels; and

wherein the means for processing the each received data packet comprises means for transferring the data in the data packet directly from the buffer to a memory associated with the output device corresponding to the indication of the first and second channels using the direct memory access engine of the channel, thereby bypassing copying of the data by the operating system to system level memory.

10. The apparatus of claim 9, further comprising:

means for sending the server a list of files which may be accessed by the client;

means for receiving from the server an indicator of a resource allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read an amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

11. The apparatus of claim 8, further comprising:

means for sending the server a list of files which may be accessed by the client;

means for receiving from the server an indicator of a resource allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read an amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

12. A network interface for controlling flow of data over a computer network from a server to one or more channels at a client, comprising:

a buffer for receiving and storing data packets from the server, wherein each data packet includes an indication of one of the channels at the client;

a controller that sends a request to a server indicating one of the channels at the client, a file to be accessed and a desired amount of data to be read from the file and that controls flow of transfer of the desired amount of data by sending a sequence of messages to a server authorizing data to be transferred from the file in an amount less than an amount of space available in the buffer on the network interface; and

means for processing the received data to provide the received data from the buffer to a de-

vice associated with the indicated channel.

13. The network interface of claim 12, wherein the means for processing comprises:

a first channel having a direct memory access engine for accessing memory in a video output device;

a second channel having a direct memory access engine for accessing memory in an audio output device;

wherein the indication of the channel at the client includes an indication of one of the first and second channels; and

means for transferring the data in each data packet directly from the buffer to the memory of the output device corresponding to the indication of the first and second channels using the direct memory access engine of the channel, thereby bypassing copying of the data by an operating system to a system level memory.

14. A process for controlling flow of data over a computer network from a server to a client, comprising:

receiving a request from the client indicating a channel at the client, a file to be accessed and a desired amount of data to be read from the file;

retrieving data from the file to be accessed and storing the data in a buffer;

receiving a message from the client authorizing an amount of data to be transferred from the file;

sending data packets to the client, wherein each data packet includes an indication of the channel at the client and data read from the buffer; and

repeating steps of receiving the message and sending data packets until the amount of data to be read from the file is received.

15. An apparatus for controlling flow of data over a computer network from a server to a client, comprising:

means for receiving a request from the client indicating a channel at the client, a file to be accessed and a desired amount of data to be read from the file;

means for retrieving data from the file to be accessed and storing the data in a buffer;

means for receiving a message from the client authorizing an amount of data to be transferred from the file;

means for sending data packets to the client, wherein each data packet includes an indication of the channel at the client and data read

from the buffer; and

wherein the means for receiving the message and sending data packets cooperate to transfer to the client the desired amount of data.

16. A process for transferring media data for a plurality of streams from a server over a computer network to a client during synchronized playback of the media data, wherein the client has one or more output devices for playing back each stream, and wherein each stream has a buffer for storing media data before playback, the process comprising:

determining an amount of data in the buffer for the output device for each stream;

selecting one of the streams, wherein the selected stream has the least amount of data in the buffer for the output device for the stream, and selecting a desired amount of data to be read for the stream;

for the selected stream, issuing one or more requests for data from the server over the computer network until the desired amount of data to be read is received from the server;

processing the data received from the server to provide the received data from the computer network to the buffer for the output device associated with the selected stream for playback by the output device; and

repeating steps of determining, selecting, issuing and processing during synchronized playback of the media data of the plurality of streams.

17. The process of claim 16, wherein issuing requests comprises:

sending a request to the server indicating the selected stream at the client, a file to be accessed and the desired amount of data to be read from the file;

sending a message to a server authorizing data to be transferred from the file in an amount less than an amount of space available in a buffer on the network interface; and

receiving data packets from the server and storing the data packets into the buffer, wherein each data packet includes an indication of the selected stream at the client.

18. The process of claim 17, wherein one of the output devices is a video output device and another of the output devices is an audio output device, the process further comprising:

associating the video output device with a first one of the streams, wherein the first stream has

a direct memory access engine;
 associating the audio output device with a second one of the streams, wherein the second stream has a direct memory access engine;

wherein the indication of the stream includes an indication of one of the first and second stream; end

wherein processing the data received from the server comprises transferring the data directly from a buffer of the network interface to the buffer associated with the output device corresponding to the indication of the first and second streams using the direct memory access engine for the stream, thereby bypassing copying of the data by the operating system to system level memory.

19. The process of claim 18, further comprising:

sending the server a list of files which may be accessed by the client; 20
 receiving from the server an indicator of a resource allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read a desired amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

20. The process of claim 16, wherein one of the output devices is a video output device and another of the output devices is an audio output device, the process further comprising:

associating the video output device with a first one of the streams, wherein the first stream has a direct memory access engine; 30
 associating the audio output device with a second one of the streams, wherein the second stream has a direct memory access engine;

wherein the indication of the stream includes an indication of one of the first and second streams; and

wherein processing the data received from the server comprises transferring the data directly from a buffer of the network interface to the buffer associated with the output device corresponding to the indication of the first and second streams using the direct memory access engine of the stream, thereby bypassing copying of the data by the operating system to system level memory.

21. The process of claim 20, further comprising:

sending the server a list of files which may be accessed by the client;

receiving from the server an indicator of a resource allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read a desired amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

22. The process of claim 16, further comprising:

sending the server a list of files which may be accessed by the client; 15
 receiving from the server an indicator of a resource allocated at the server for storing the data of each of the files; and

wherein the request sent to the server to read a desired amount of data from a file includes the indicator of the resource allocated for storing the data of the file.

23. The process of claim 16, wherein issuing one or more requests comprises issuing requests according to an amount of space available in a buffer of the client.

24. The process of claim 16, wherein issuing one or more requests comprises issuing a request only when space is available in a buffer in the client.

25. The process of claim 16, wherein issuing one or more requests comprises issuing a request so as to ensure that the flow of data is within a capability of the client to receive data for delivery to the buffer for the output device of the selected stream.

26. The process of claim 16, wherein issuing one or more requests comprises issuing a request such that each request permits the server to send only as much data as the client can reliably receive.

27. The process of claim 16, wherein issuing one or more requests comprises issuing a request such that each request permits the server to send only as much data as the client can receive reliably at a full speed of the computer network.

28. The process of claim 16, wherein issuing one or more requests comprises issuing a request such that each request ensures that the server does not transmit data faster than the client can accept the data.

29. The process of claim 16, wherein issuing one or more requests comprises issuing a request such that each request authorizes the server to send only an amount of data which a buffer in the client is ca-

pable of receiving.

30. A system for transferring media data for a plurality of streams from a server over a computer network to a client during synchronized playback of the media data, wherein the client has one or more output devices for playing back each stream, and wherein each stream has a buffer for storing media data before playback, the process comprising:

5

10

15

20

25

means for determining an amount of data in the buffer for the output device for each stream; means for selecting one of the streams, wherein the selected stream has the least amount of data in the buffer for the output device for the stream, and selecting a desired amount of data to be read for the stream; means for issuing one or more requests for data from the server over the computer network for the selected stream until the desired amount of data to be read is received from the server; means for processing the data received from the server to provide the received data from the computer network to the buffer for the output device associated with the selected stream for playback by the output device; and

wherein the means for determining, selecting, issuing and processing cooperate during synchronized playback of the media data of the plurality of streams.

35

40

45

50

55

13




Fig. 1




Fig. 2




Fig. 3
(Prior Art)




Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8